

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
6 February 2003 (06.02.2003)

PCT

(10) International Publication Number
WO 03/010680 A1

(51) International Patent Classification⁷: **G06F 15/16**

(21) International Application Number: PCT/US02/23811

(22) International Filing Date: 25 July 2002 (25.07.2002)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
09/916,930 26 July 2001 (26.07.2001) US

(71) Applicant: **NETWORKS ASSOCIATES TECHNOLOGY, INC.** [US/US]; 3965 Freedom Circle, Santa Clara, CA 95054 (US).

(72) Inventors: **ROTHWELL, Anton, C.**; 7 Parmiter Close, Aylesbury, Buckinghamshire HP19 8GS (GB). **JAGGER,**

Luke, D.; 3 Egypt Way, Aylesbury, Buckinghamshire HP19 8GP (GB). **DENNIS, William, R.**; 3 Nuthatch, Aylesbury, Buckinghamshire HP19 0WF (GB).


(74) Agent: **ZILKA, Kevin, J.**; Silicon Valley IP Group, LLC., P.O. Box 721120, San Jose, CA 95172-1120 (US).

(81) Designated States (*national*): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZM, ZW.

(84) Designated States (*regional*): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE,

[Continued on next page]

(54) Title: INTELLIGENT SPAM DETECTION SYSTEM USING AN UPDATEABLE NEURAL ANALYSIS ENGINE



(57) Abstract: A system, method, and computer program product are provided for detecting an unwanted message. First, an electronic mail message (402) is received. Text in the electronic mail message is decomposed (406). Statistics associated with the text are gathered using a statistical analyzer (408). A neural network engine (410) coupled to the statistical analyzer is taught to recognize unwanted messages based on statistical indicators. The statistical indicators are analyzed utilizing the neural network engine for determining whether the electronic mail message is an unwanted message.

WO 03/010680 A1



ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, SK,
TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ,
GW, ML, MR, NE, SN, TD, TG).

*For two-letter codes and other abbreviations, refer to the "Guid-
ance Notes on Codes and Abbreviations" appearing at the begin-
ning of each regular issue of the PCT Gazette.*

Published:

— *with international search report*

INTELLIGENT SPAM DETECTION SYSTEM USING AN UPDATEABLE NEURAL ANALYSIS ENGINE

5

FIELD OF THE INVENTION

The present invention relates to SPAM detection methods, and more particularly to intelligently detecting and removing SPAM.

10

BACKGROUND OF THE INVENTION

The rapid increase in the number of users of electronic mail and the low cost of distributing electronic messages, for example, via the Internet and other communications networks has made mass marketing via electronic mail ("e-mail") an attractive advertising medium. Consequently, e-mail is now frequently used as the medium for widespread marketing broadcasts of unsolicited messages to e-mail addresses, commonly known as "SPAM."

20

Electronic mass marketers (also called "spammers") use a variety of techniques for obtaining e-mail address lists. For example, marketers obtain e-mail addresses from postings on various Internet sites such as news group sites, chat room sites, or directory services sites, message board sites, mailing lists, and by identifying "mailto" address links provided on web pages. Using these and other similar methods, electronic mass marketers may effectively obtain large numbers of mailing addresses, which become targets for their advertisements and other unsolicited messages.

25

Users of Internet services and electronic mail, however, are not eager to have their e-mail boxes filled with unsolicited e-mails. This is an increasing problem for Internet service providers (ISPs) such as America Online (AOL®) or Microsoft Network (MSN®) and other entities with easily identifiable e-mail addresses such as large corporations (e.g., IBM®, Microsoft®, General Motors®, etc.). ISPs object to junk mail because it reduces their users' satisfaction of their services. Corporations want to eliminate junk mail because it reduces worker productivity.

To date, the prior art has been devoid of mechanisms that can block SPAM effectively. Traditionally, SPAM detection has been based around specific rules for detecting it. Such rules include searching for key phrases in the subject headers, determining whether the recipient is actually on the list of users to receive the e-mail, etc.

More particularly, prior art systems rely on a set rule-base or on blocking based on simple mail fields such as sender, subject, mail body, etc. As Spammers become more creative in their mailings, it is increasingly difficult to block unwanted messages based on fixed information. Text search mechanisms in particular often generate a number of misses or "falses" due to the limitations of the searching mechanisms.

In particular, such text search mechanisms traditionally utilize static logic for locating particular known strings. This is often insufficient due to the dynamic manner in which spamming methods change over time. Thus, what is needed is a process for dynamically and intelligently detecting unwanted SPAM electronic mail messages.

DISCLOSURE OF THE INVENTION

A system, method and computer program product are provided for detecting an unwanted message. First, an electronic mail message is received. Text in the
5 electronic mail message is decomposed. Statistics associated with the text are gathered using a statistical analyzer. A neural network engine coupled to the statistical analyzer is taught to recognize unwanted messages based on statistical indicators. The statistical indicators are analyzed utilizing the neural network engine for determining whether the electronic mail message is an unwanted message.

10

As mentioned above, the neural network engine can be taught to recognize unwanted messages. In one process of teaching the neural network, examples are provided to the neural network engine. The examples are of wanted messages and unwanted messages. Each of the examples is associated with a desired output. Each
15 of the examples is processed with statistics by the neural network engine for generating weights for the statistics. Each of the weights is used to denote wanted and unwanted messages. Preferably, the neural network engine utilizes adaptive linear combination for adjusting the weights. Logic associated with the neural network engine is updated based on the processing by the neural network engine.

20

In another process for teaching the neural network engine, the neural network engine is updated to recognize an unwanted message. The message is identified as an unwanted message. The features of the message that make the message unwanted are identified, wherein the identified features are stored and used by the neural
25 network to identify subsequent unwanted messages. Preferably, a graphical user interface is provided for allowing a user to identify the features of the message that make the message unwanted.

In another aspect of the present embodiment, the neural network engine uses artificial intelligence that analyzes previous user input for determining whether the message is unwanted.

- 5 A system, method and computer program product are also provided for teaching a neural network engine to recognize an unwanted message. Examples are provided to a neural network engine. The examples are of wanted messages and unwanted messages. Each of the examples is associated with a desired output. Each of the examples is processed with statistics for generating weights for the statistics.
- 10 Each of the weights is used to denote wanted and unwanted messages. Logic associated with the neural network engine is updated based on the processing by the neural network engine.

- In one aspect of the present embodiment, the neural network engine utilizes
- 15 adaptive linear combination for adjusting the weights. In another aspect of the present embodiment, the neural network engine is updated to recognize an unwanted message. The message is identified as an unwanted message. The features of the message that make the message unwanted are identified. The identified features are stored and used by the neural network to identify subsequent unwanted messages.
- 20 Preferably, a graphical user interface is provided for allowing a user to identify the features of the message that make the message unwanted.

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 illustrates a network architecture in accordance with the one
5 embodiment.

Figure 2 shows a representative hardware environment associated with the
computers of Figure 1.

10 Figure 3 is a flow diagram of a process for detecting an unwanted message.

Figure 4 depicts an illustrative architecture according to an embodiment.

Figure 5 is a flowchart of a process for teaching a neural network engine to
15 recognize an unwanted message.

Figure 6 is a flow diagram depicting processing performed by the neural
network engine.

20 Figure 7 illustrates a system for allowing a user to teach the neural network
engine to recognize unwanted messages.

Figure 8 illustrates a user interface used in conjunction with the system of
figure 7.

25

DESCRIPTION OF THE PREFERRED EMBODIMENTS

Figure 1 illustrates a network architecture 100, in accordance with the one
5 embodiment. As shown, computers 102 of remote users are connected to a network
104. The remote users send electronic mail messages (e-mail) to local users, who
receive them on computers 106. In the context of the present network architecture,
the network may take any form including, but not limited to a local area network
(LAN), a wide area network (WAN) such as the Internet, etc. The computers can
10 include a desktop computer, laptop computer, hand-held computer, etc.

The e-mail passes through a gateway 108 which analyzes the messages to
determine whether they are SPAM prior to allowing the message to pass to the local
users. In one embodiment, the gateway 108 may include an Internet gateway,
15 intranet gateway, Internet proxy, intranet proxy, or any other type of intermediate
device. The gateway includes a statistical analyzer 110 and a neural network engine
112.

In use, the statistical analyzer decomposes the electronic messages to
20 determine an amount of various SPAM indicators (i.e. capitalization, punctuation,
URLs, phone numbers, etc.). Instead of using static rule-based logic to accept or
reject the electronic messages based on the results of the statistical analyzer, the
results of the parsing are passed to the neural network engine. The neural network
engine is first "taught" the difference between SPAM and non-SPAM electronic
25 messages based on the aforementioned SPAM indicators. This is done by
processing hundreds of electronic messages that are known to be SPAM or known to
be non-SPAM. Once taught, the neural network engine can be used in combination
with the statistical analyzer to accept or deny electronic messages. An administrator
114 in communication with the gateway can be allowed to manipulate operation of
30 the gateway and its components.

Figure 2 shows a representative hardware environment that may be associated with the remote source 102 and/or target 106 of Figure 1, in accordance with one embodiment. Such figure illustrates a typical hardware configuration of a workstation in accordance with a preferred embodiment having a central processing unit 210, such as a microprocessor, and a number of other units interconnected via a system bus 212.

The workstation shown in Figure 2 includes a Random Access Memory (RAM) 214, Read Only Memory (ROM) 216, an I/O adapter 218 for connecting peripheral devices such as disk storage units 220 to the bus 212, a user interface adapter 222 for connecting a keyboard 224, a mouse 226, a speaker 228, a microphone 232, and/or other user interface devices such as a touch screen (not shown) to the bus 212, communication adapter 234 for connecting the workstation to a communication network 235 (e.g., a data processing network) and a display adapter 236 for connecting the bus 212 to a display device 238.

The workstation may have resident thereon an operating system such as the Microsoft Windows NT or Windows/95 Operating System (OS), the IBM OS/2 operating system, the MAC OS, Linux or other UNIX operating system. It will be appreciated that a preferred embodiment may also be implemented on platforms and operating systems other than those mentioned. A preferred embodiment may be written using JAVA, C, and/or C++ language, or other programming languages, along with an object oriented programming methodology. Object oriented programming (OOP) has become increasingly used to develop complex applications.

SPAM DETECTION

Figure 3 is a flow diagram of a process 300 for detecting an unwanted message. In operation 302, an electronic mail message is received. Text in the electronic mail message is decomposed in operation 304. Statistics associated with

the text are gathered in operation **306** using a statistical analyzer. This may be accomplished utilizing various methods such as analyzing character type, i.e. Unicode, etc. In operation **308**, a neural network engine coupled to the statistical analyzer is taught to recognize unwanted messages based on statistical indicators.

5 The statistics are analyzed in operation **310** utilizing the neural network engine for determining whether the electronic mail message is an unwanted message.

According to another embodiment, a method is provided by which it is possible to accurately detect SPAM by applying statistical calculations against the

10 text content of the message and supply the results of the application of the calculations and the message text contents itself into a neural network engine. The neural network then attempts to determine whether the message is SPAM or not based on what the neural network has learned in the past and/or by comparison with an existing set of known SPAM. An additional mechanism can be provided so that a

15 user can return a message to the SPAM engine and mark it as SPAM (or not SPAM) to provide the engine with an on-going learning capability.

ARCHITECTURE

20 Figure **4** depicts an illustrative gateway architecture **400** according to an embodiment. The e-mail **402** arrives from the Internet **404** at the word decomposer **406**, which breaks the text content into words and punctuation. The parsed text is then supplied to the statistical analyzer **408** which creates a table of variables relating to the message, e.g.: Total Number of words, Number of words capitalized,

25 Punctuation to word ratio etc. See Table **1** and related discussion, below.

This table along with the decomposed word list is supplied to the neural network **410** which provides a weighting, or probability, that the message is SPAM, partially based on known patterns of SPAM messages stored in a pattern database

30 **412**. If the message is determined to be SPAM, it is quarantined in a quarantine

database 414. If the message is determined not to be SPAM, the message is sent to a mail forwarder 416, which forwards the message to a user 418.

STATISTICAL WORD ANALYZER

5

The Statistical Word Analyzer attempts to build some key facts about the text content. The facts are based on certain characteristics that users/administrators have determined could represent SPAM. A results table from analysis of the message is then built.

10

At that point, the program segment shown in Example 1 may be used.

Example 1

15

```
If words_capitalized >5% and total_punc>20%  
  And perc_plink >50% And URL analysis fails  
  And E-mail address analysis fails Then  
    This is spam
```

```
Else
```

20

```
  This is not Spam  
End If
```

25

However, questions may arise as to whether the above analysis is accurate for all SPAM, whether the percentages are right, or whether the test variables are the correct ones to use.

30

Because this task is difficult to do using a fixed algorithm, the statistical analysis ends at this point and this data is passed to the Neural Network engine to determine patterns in statistics and words, and use these to determine whether the message is SPAM based on comparing the patterns to patterns predetermined to be SPAM or non-SPAM. The greater the number of variables in the statistics table, the easier it is for the Artificial Intelligence engine (AI) to “learn” to differentiate between SPAM and genuine messages.

The AI solution provides two goals. In particular, the AI is used to produce a set of rules that can be used in an existing AI engine. Further, the AI engine is used as a standalone gateway for determining which messages are SPAM.

5

Table 1 illustrates various steps and/or functions to achieving these goals.

Table 1

10	Pre-train the system with known SPAM from an archive of known SPAM, such as http://www.annexia.org/spam/ "The Great Spam Archive".
15	User Interface (UI) modification of the AI. A point and click UI where an existing e-mail is selected and displayed to be made an example of. Rules are constructed from the areas of the example mail that classify it as SPAM. See the section entitled Feature Editor Applet, below for a more detailed description.
20	Application of a tiered approach, in which the user sends the message to an administrative area. An administrator utilizes the UI to notify the engine.
25	Artificial Intelligence introduced to make decisions based on previous Administrative input. This may include statistical or pattern-matching intelligence and would automatically update the rule-base.
30	Artificial Intelligence taken a level further, where grammatical and language decisions are made based on previous human-input to provide automatic generation of a rule-base.
35	System opened up to internal users with encryption method for trusted SPAM identification by the users.

NEURAL NETWORK ENGINE

The statistics table is passed as inputs to the Artificial Neural Network (ANN). The preferred ANN is a Supervised Learning type, though other types can be used. In this type, a “teacher” (user, administrator, or computer application) shows examples of inputs that the engine will receive paired with desired outputs. An error value is produced between the desired and actual responses, which should reduce as the learning progresses.

Figure 5 is a flowchart of a process 500 for teaching a neural network engine to recognize an unwanted message. In operation 502, examples are provided to a neural network engine. The examples are of wanted messages and unwanted messages. Each of the examples is associated with a desired output. In operation 504, each of the examples is processed with statistics for generating weights for the statistics. Each of the weights is used to denote wanted and unwanted messages. Logic associated with the neural network engine is updated in operation 506 based on the processing by the neural network engine.

In the SPAM context, there are two sets of inputs: First, an archive containing only SPAM is inputted, and secondly an archive containing only genuine (non-SPAM) messages is inputted. Known SPAM can be obtained from various online resources (<http://www.annexia.org/spam/> “The Great Spam Archive”). The teacher may automatically and randomly pick messages from either archive and supply them (with the statistical table) to the ANN together with a value for SPAM or non-SPAM. This continues until the archives are exhausted. A very large data set possibly consisting of over 500,000 messages is preferred.

Each input message is expressed as a vector, each variable from the statistic table being a feature variable in the vector. Note Table 2.

Table 2

		Capitals
	X=	Punctuation
		Dollars
5	...	

Figure 6 is a flow diagram depicting processing performed by the ANN. The most appropriate form of ANN would be an Adaptive Linear Combiner (ALC) 600, which allows the presentation of input vectors X 602 and desired responses d 604 to the ALC. This will adjust weights until outputs a 606 are close to the desired responses.

After the learning process has taken place, the Engine can be deployed into the gateway situation. All associated vectors, matrices and weights to be used with the ALC can be stored permanently on disk.

The gateway could also be used to monitor intercommunication between Internet servers for tighter company-wide security, such as for preventing outbreak, SPAM, hacking attempts, etc. Such functionality can extend to the whole Internet community.

In addition to the pre-learning, there can also be an interactive learning mechanism while the gateway is active. This is discussed in the next section.

25 USER TEACHING MECHANISM

Some users of the system will be given permissions to allow teaching the Spam Engine when new mail arrives. Figure 7 illustrates the system that allows such teaching.

A typical sequence of actions using the system would be as follows. E-mail 702 is received by a user teacher 704, but the e-mail has been marked incorrectly as

clean. The user returns the e-mail 706 back to the Gateway 708 but indicates the desired response as "Spam." A learner 710 in the Gateway receives the e-mail and desired response. The Gateway stores the e-mail creating a unique ID for it. The Gateway sends a new e-mail 712 to the user. The new e-mail contains a unique ID
5 contained within a URL to the Feature Editor (Java Applet) 714. The user clicks on URL and indicates which features of the text of the e-mail make it SPAM. Preferably, the user is allowed to select the parts of the mail that clearly define it to be SPAM/offensive (e.g. subject, body, sender, attachments etc) and then within each part refine it down to the data elements to search on for CF filtering (set of
10 words, file name / type for attachment, sender domain etc). The results of the user input are passed into a feature queue 716.

The new features are not applied directly to the Statistical Analyzer/AI Engine 718 in one embodiment because the ANN would have to be reset and re-learn its
15 mail archive 720. Instead, a batch teacher 722 schedules an entire learning run at a scheduled interval with additional features from the feature queue. A secondary ANN 724 is used so that the primary ANN can stay active and continue to process e-mail while the secondary ANN is learning. When the learning process is finished, the primary ANN becomes the secondary ANN and vice versa. The AI Engine is
20 then initialized with the new features.

The AI engine could optionally be configured to divide SPAM into different confidence levels, e.g. Definitely SPAM, definitely not SPAM and possibly SPAM. The possibly SPAM would still be forwarded to the recipient, but also copied to a
25 user teacher for feeding manual categorization in future training.

FEATURE EDITOR APPLET

Figure 8 illustrates the user interface 800 for applying features to the text of
30 an e-mail. The user is able to highlight areas of text and select (right click) from a

pop up menu 802 a measurement to apply to this text. Such measurements are shown in Table 3.

Table 3

5

Contains a single instance of word = true/false
Contains multiple instances of word = true/false
Count number of occurrences of word = Number
Ratio of word to total number of words = Percentage

10

Selection of multiple words is also possible. Additional measurements will then be applied. Note Table 4.

Table 4

15

All words must exist = true/false
Some words can exist = true/false
Multiple instances of all words must exist = true/false
Multiple instances of some words must exist = true/false
Ratio of all words to total number of words = Percentage
Ratio of some words to total number of words = Percentage

20

When the OK button 804 is selected, these additional measurements are added to the feature queue for future inclusion in a batch teaching session. It is also possible for an administrator to review and edit the additional features before a batch run takes place.

25

While various embodiments have been described above, it should be understood that they have been presented by way of example only, and not limitation. Thus, the breadth and scope of a preferred embodiment should not be limited by any of the above-described exemplary embodiments, but should be

30

defined only in accordance with the following claims and their equivalents.

CLAIMS

What is claimed is:

1

1. A method for detecting an unwanted message, comprising:
 - (a) receiving an electronic mail message;
 - (b) decomposing text in the electronic mail message;
 - (c) gathering statistics associated with the text using a statistical analyzer;
 - (d) teaching a neural network engine coupled to the statistical analyzer to recognize unwanted messages based on statistical indicators; and
 - (e) analyzing the statistical indicators utilizing the neural network engine for determining whether the electronic mail message is an unwanted message.
2. The method as recited in claim 1, wherein the neural network engine is taught to recognize unwanted messages by providing examples to the neural network engine, wherein the examples are of wanted messages and unwanted messages, wherein each of the examples is associated with a desired output.
3. The method as recited in claim 2, wherein each of the examples are processed with statistics by the neural network engine for generating weights for the statistics, wherein each of the weights is used to denote wanted and unwanted messages.
4. The method as recited in claim 3, wherein the neural network engine utilizes adaptive linear combination for adjusting the weights.
5. The method as recited in claim 3, wherein logic associated with the neural network engine is updated based on the processing by the neural network engine.

6. The method as recited in claim 1, wherein the teaching of the neural network engine includes identifying a message as an unwanted message, the features of the message that make the message unwanted are identified, and the identified features are stored and used by the neural network to identify subsequent unwanted messages.
7. The method as recited in claim 6, wherein a graphical user interface is provided for allowing a user to identify the features of the message that make the message unwanted.
8. The method as recited in claim 1, wherein the neural network engine analyzes previous user input for determining whether the message is unwanted.
9. The method as recited in claim 1, wherein the statistics gathered using the statistical analyzer include results of an analysis of character type.
10. A computer program product for detecting an unwanted message, comprising:
 - (a) computer code for receiving an electronic mail message;
 - (b) computer code for decomposing text in the electronic mail message;
 - (c) computer code for gathering statistics associated with the text using a statistical analyzer;
 - (d) computer code for teaching a neural network engine coupled to the statistical analyzer to recognize unwanted messages based on statistical indicators; and
 - (e) computer code for analyzing the statistical indicators utilizing the neural network engine for determining whether the electronic mail message is an unwanted message.
11. A system for detecting an unwanted message, comprising:

- (a) a statistical analyzer for gathering statistics associated with text retrieved from an electronic mail message; and
- (b) a neural network engine coupled to the statistical analyzer, the neural network engine capable of being taught to recognize unwanted messages based on statistical indicators;
- (c) wherein the neural network engine analyzes the statistical indicators for determining whether the electronic mail message is an unwanted message.

12. A method for teaching a neural network engine to recognize an unwanted message, comprising:

- (a) providing examples to a neural network engine, wherein the examples are of wanted messages and unwanted messages, and each of the examples is associated with a desired output;
- (b) processing each of the examples with statistics for generating weights for the statistics, wherein each of the weights is used to denote wanted and unwanted messages; and
- (c) updating logic associated with the neural network engine based on the processing by the neural network engine.


13. The method as recited in claim 12, wherein the neural network engine utilizes adaptive linear combination for adjusting the weights.

14. The method as recited in claim 12, wherein the neural network engine is updated to recognize an unwanted message, the message is identified as an unwanted message, the features of the message that make the message unwanted are identified, and the identified features are stored and used by the neural network to identify subsequent unwanted messages.

15. The method as recited in claim 14, wherein a graphical user interface is provided for allowing a user to identify the features of the message that make the message unwanted.
16. A computer program product for teaching a neural network engine to recognize an unwanted message, comprising:
- (a) computer code for providing examples to a neural network engine, wherein the examples are of wanted messages and unwanted messages, and each of the examples is associated with a desired output;
 - (b) computer code for processing each of the examples with statistics for generating weights for the statistics, wherein each of the weights is used to denote wanted and unwanted messages; and
 - (c) computer code for updating logic associated with the neural network engine based on the processing by the neural network engine.
17. A system for teaching a neural network engine to recognize an unwanted message, comprising:
- (a) a neural network engine adapted to receive examples of wanted messages and unwanted messages, each of the examples being associated with a desired output, the neural network engine further adapted for processing each of the examples with statistics for generating weights for the statistics, each of the weights being used to denote wanted and unwanted messages; and
 - (b) wherein the neural network engine is capable of being updated.
18. A method for detecting an unwanted electronic mail message, comprising:
- (a) receiving user input via a graphical user interface, the user input including indicators of unwanted electronic mail messages, wherein the indicator for a

- particular unwanted message includes at least one of a URL, a sender, and domain of a sender;
- (b) storing the user input indicators in a database;
 - (c) receiving an electronic mail message prior to delivery of the electronic mail message to an intended recipient of the electronic mail message;
 - (d) analyzing the electronic mail message for determining whether the electronic mail message is an unwanted message based on the user input indicators stored in the database;
 - (e) sending the electronic mail message to the intended recipient if the electronic mail message is not determined to be unwanted; and
 - (f) selecting a disposition of the electronic mail message if the electronic mail message is determined to be unwanted, the disposition including at least one of: not delivering the electronic mail message to the intended recipient, sending the electronic mail message to an administrator, and sending the electronic mail message to a quarantine.
19. The method as recited in claim 18, and further comprising decomposing text in the electronic mail message, gathering statistics associated with the text using a statistical analyzer, and analyzing the statistical indicators utilizing a neural network engine for determining whether the electronic mail message is an unwanted message.
20. A method for detecting an unwanted message, comprising:
- (a) receiving user input via a graphical user interface, the user input including indicators of unwanted electronic mail messages;
 - (b) teaching a neural network engine to recognize unwanted messages based on the user input;
 - (c) receiving an electronic mail message; and

- (d) analyzing the electronic mail message utilizing the neural network engine for determining whether the electronic mail message is an unwanted message based on the user input;
 - (e) sending the electronic mail message to the intended recipient if the electronic mail message is not determined to be unwanted; and
 - (f) selecting a disposition of the electronic mail message if the electronic mail message is determined to be unwanted, the disposition including at least one of: not delivering the electronic mail message to the intended recipient, sending the electronic mail message to an administrator, and sending the electronic mail message to a quarantine.
21. The method as recited in claim 20, and further comprising decomposing text in the electronic mail messages, gathering statistics associated with the text using a statistical analyzer, and analyzing the statistical indicators utilizing the neural network engine for determining whether the electronic mail message is an unwanted message.



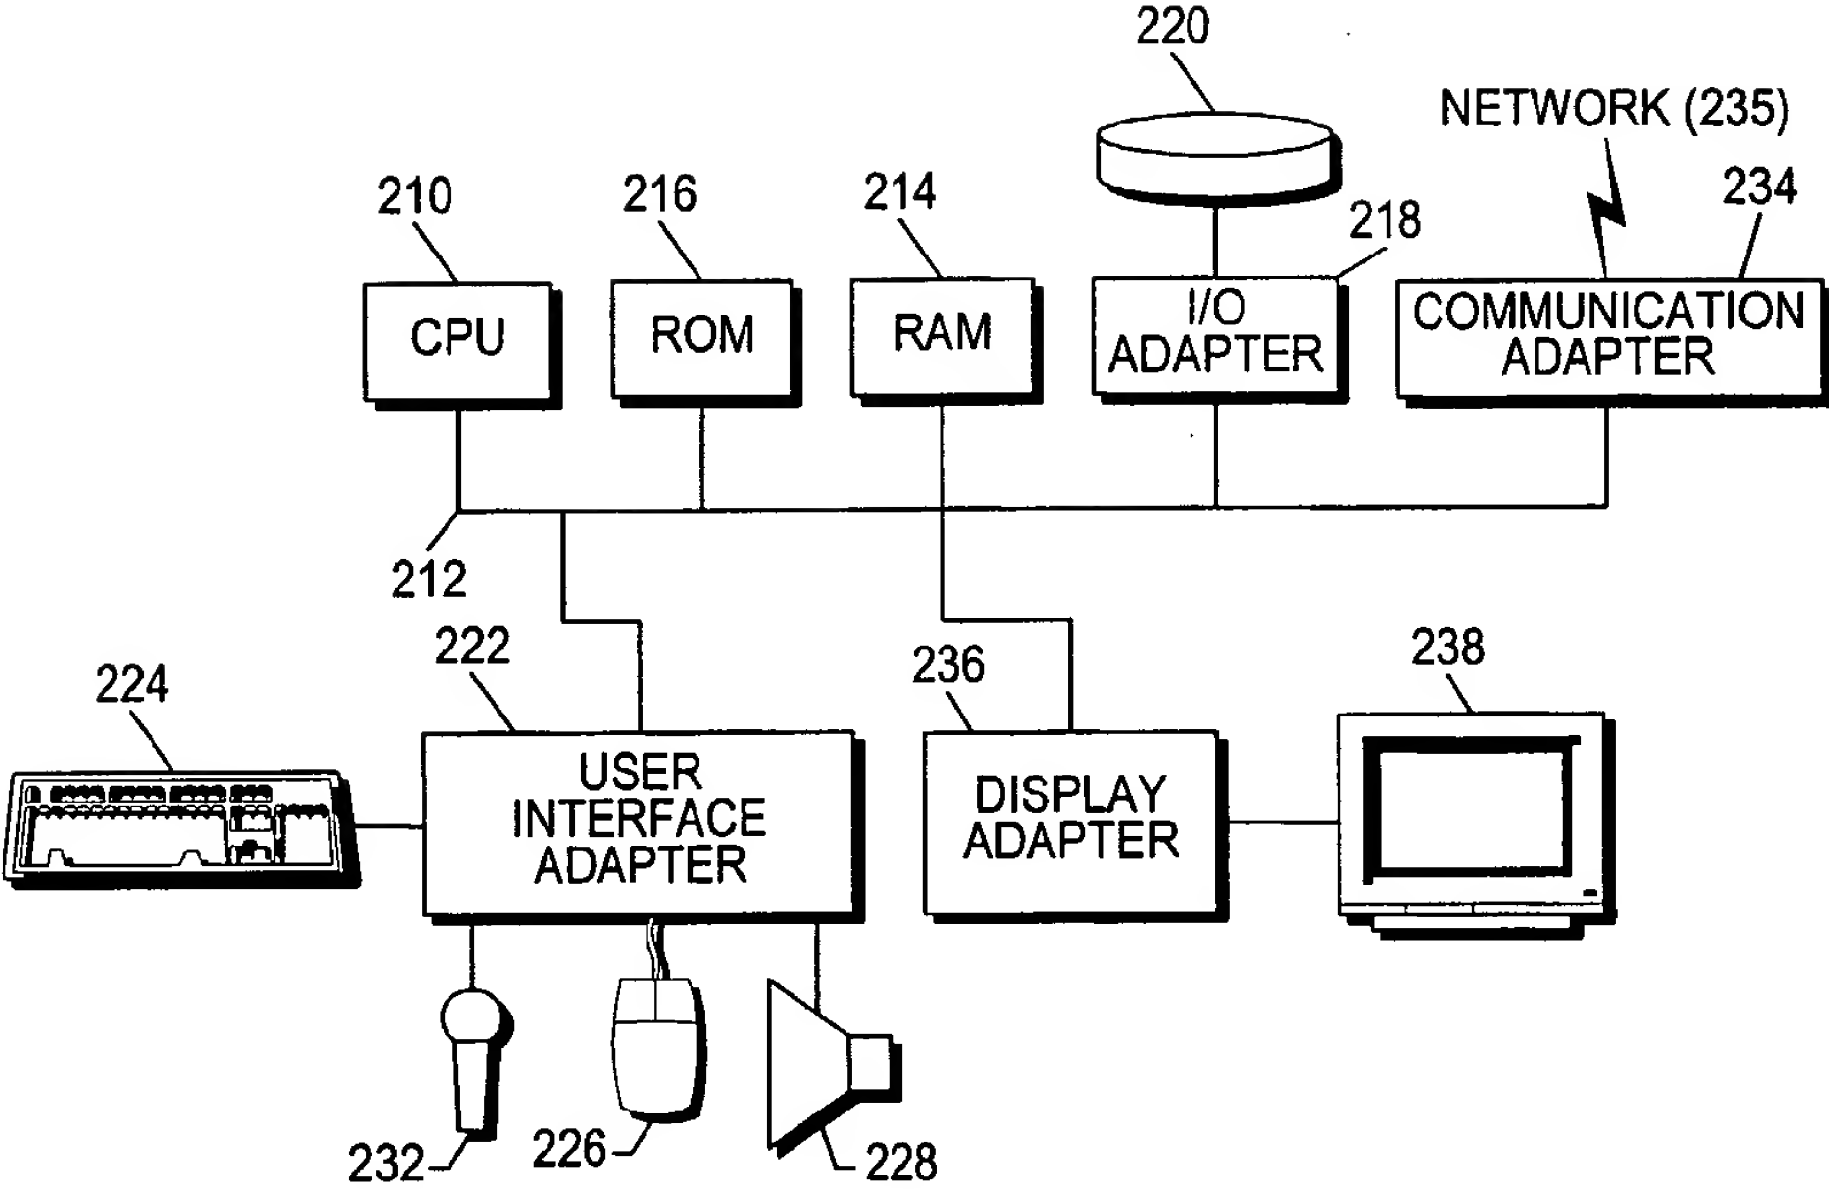


Fig. 2

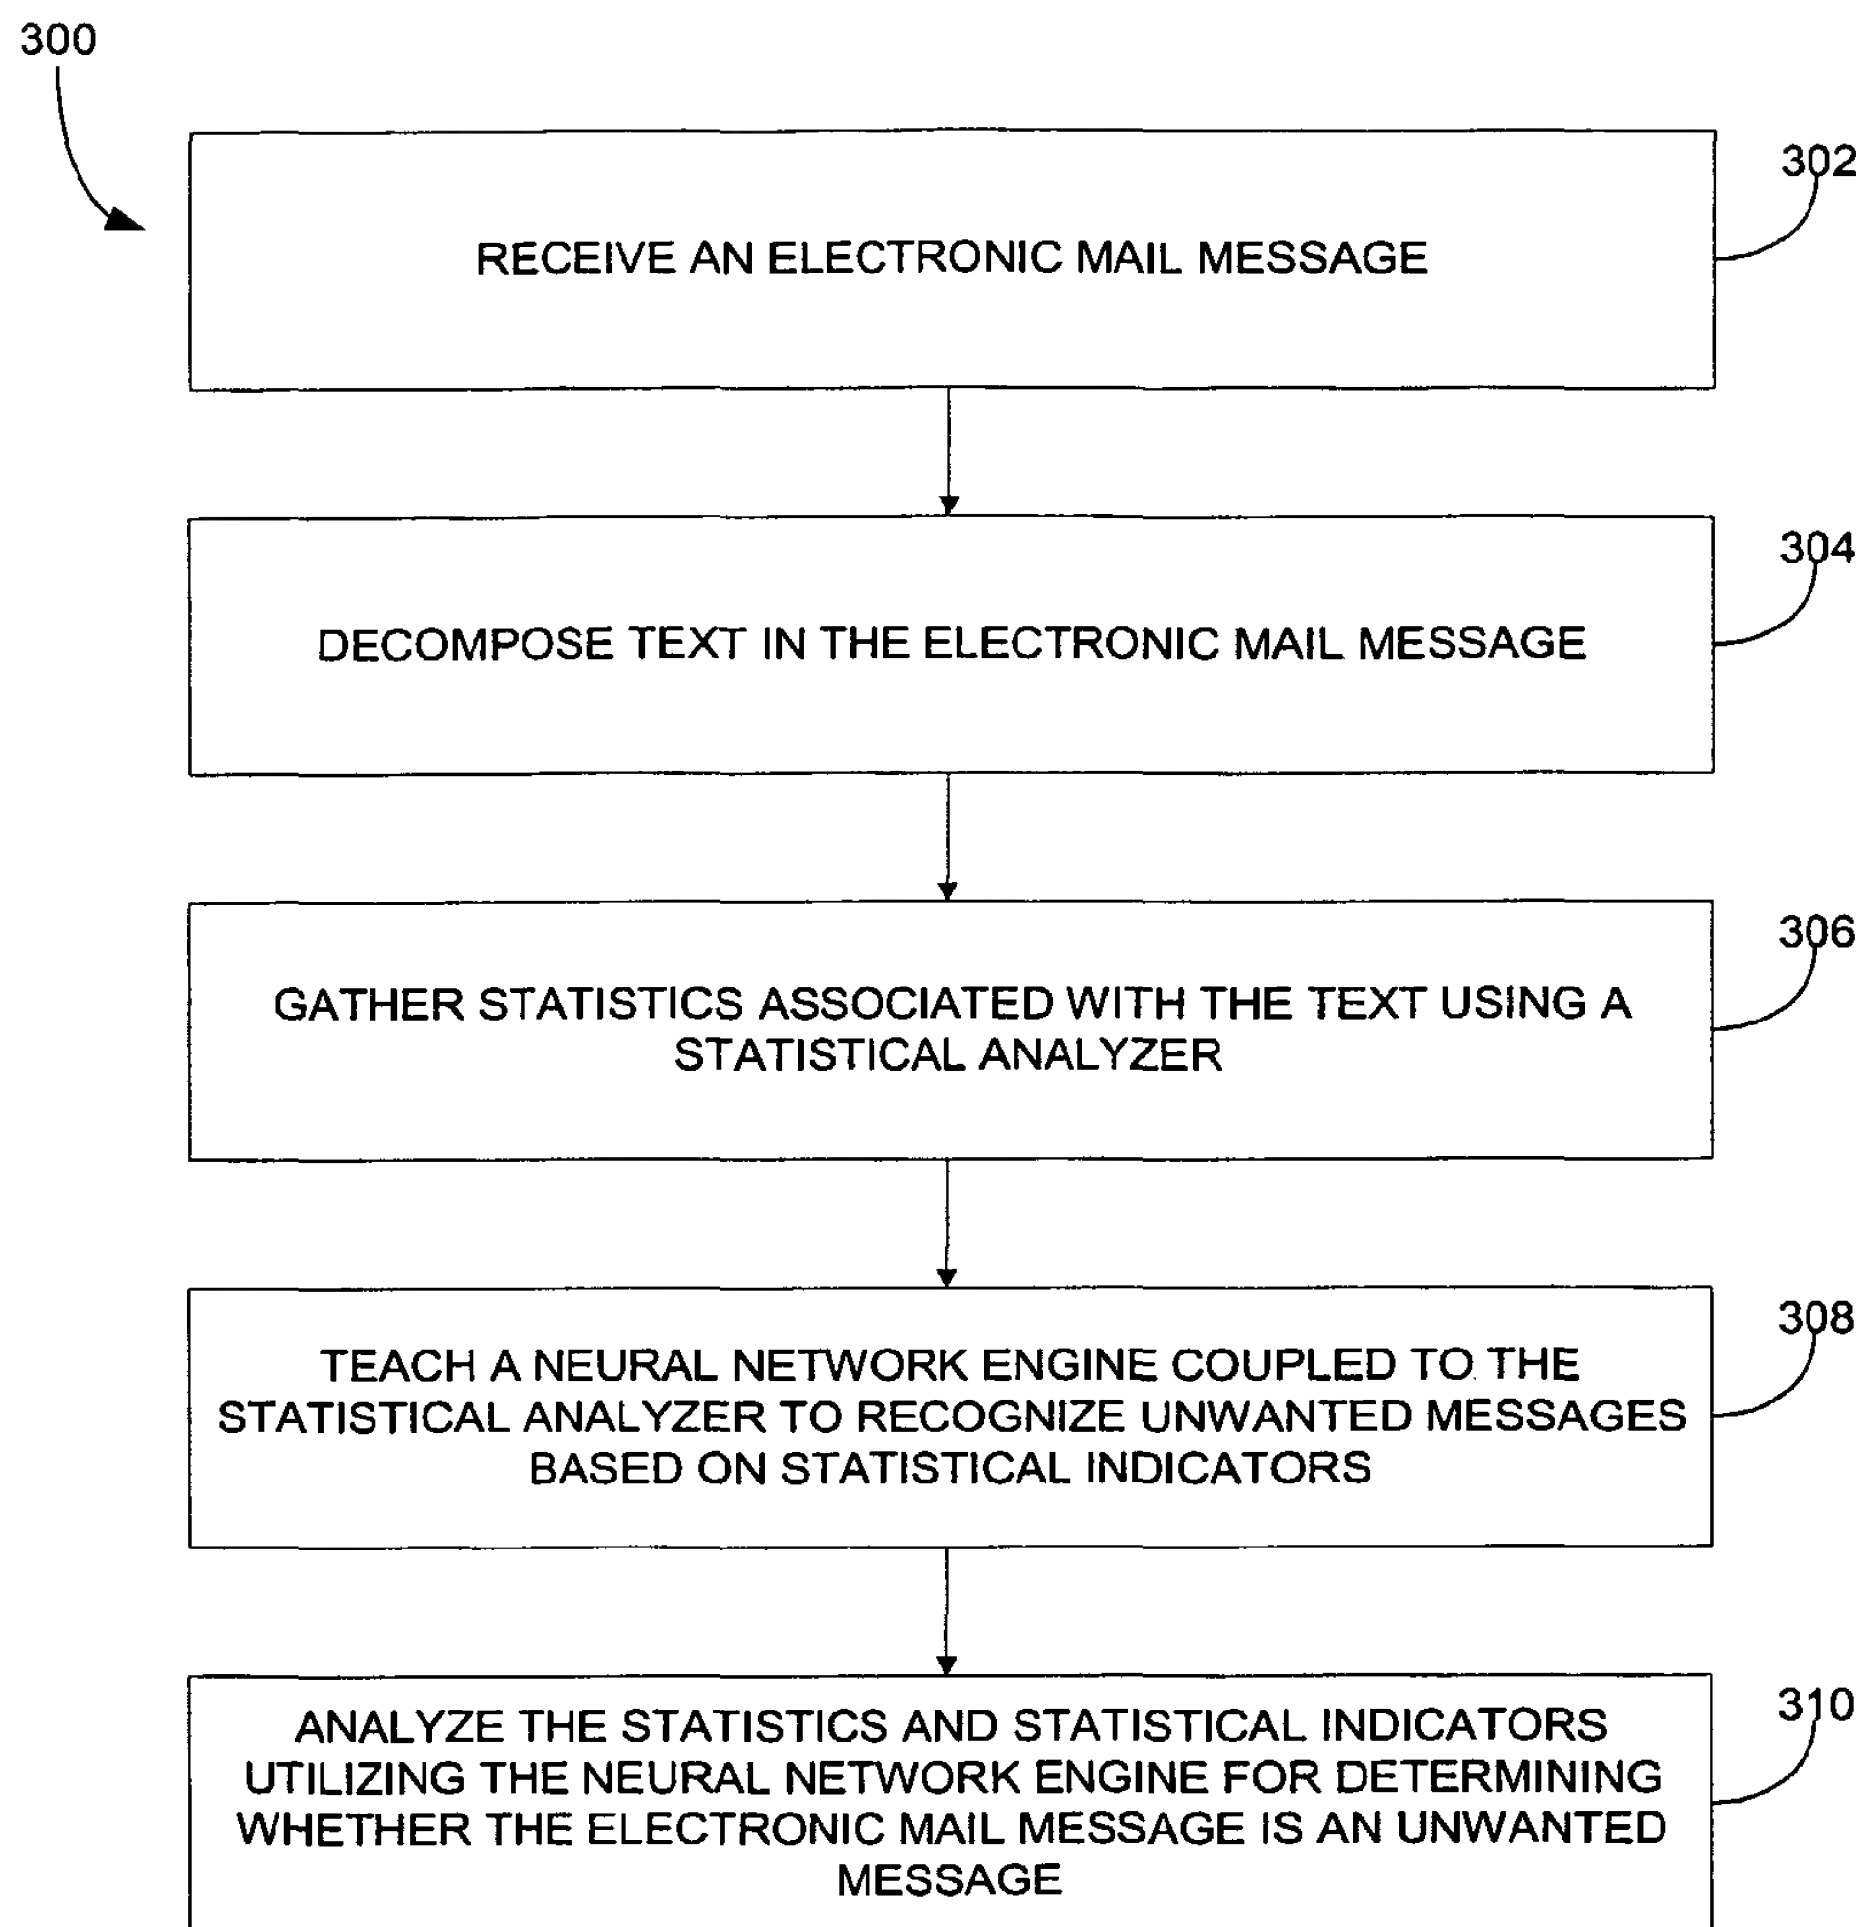


Fig. 3

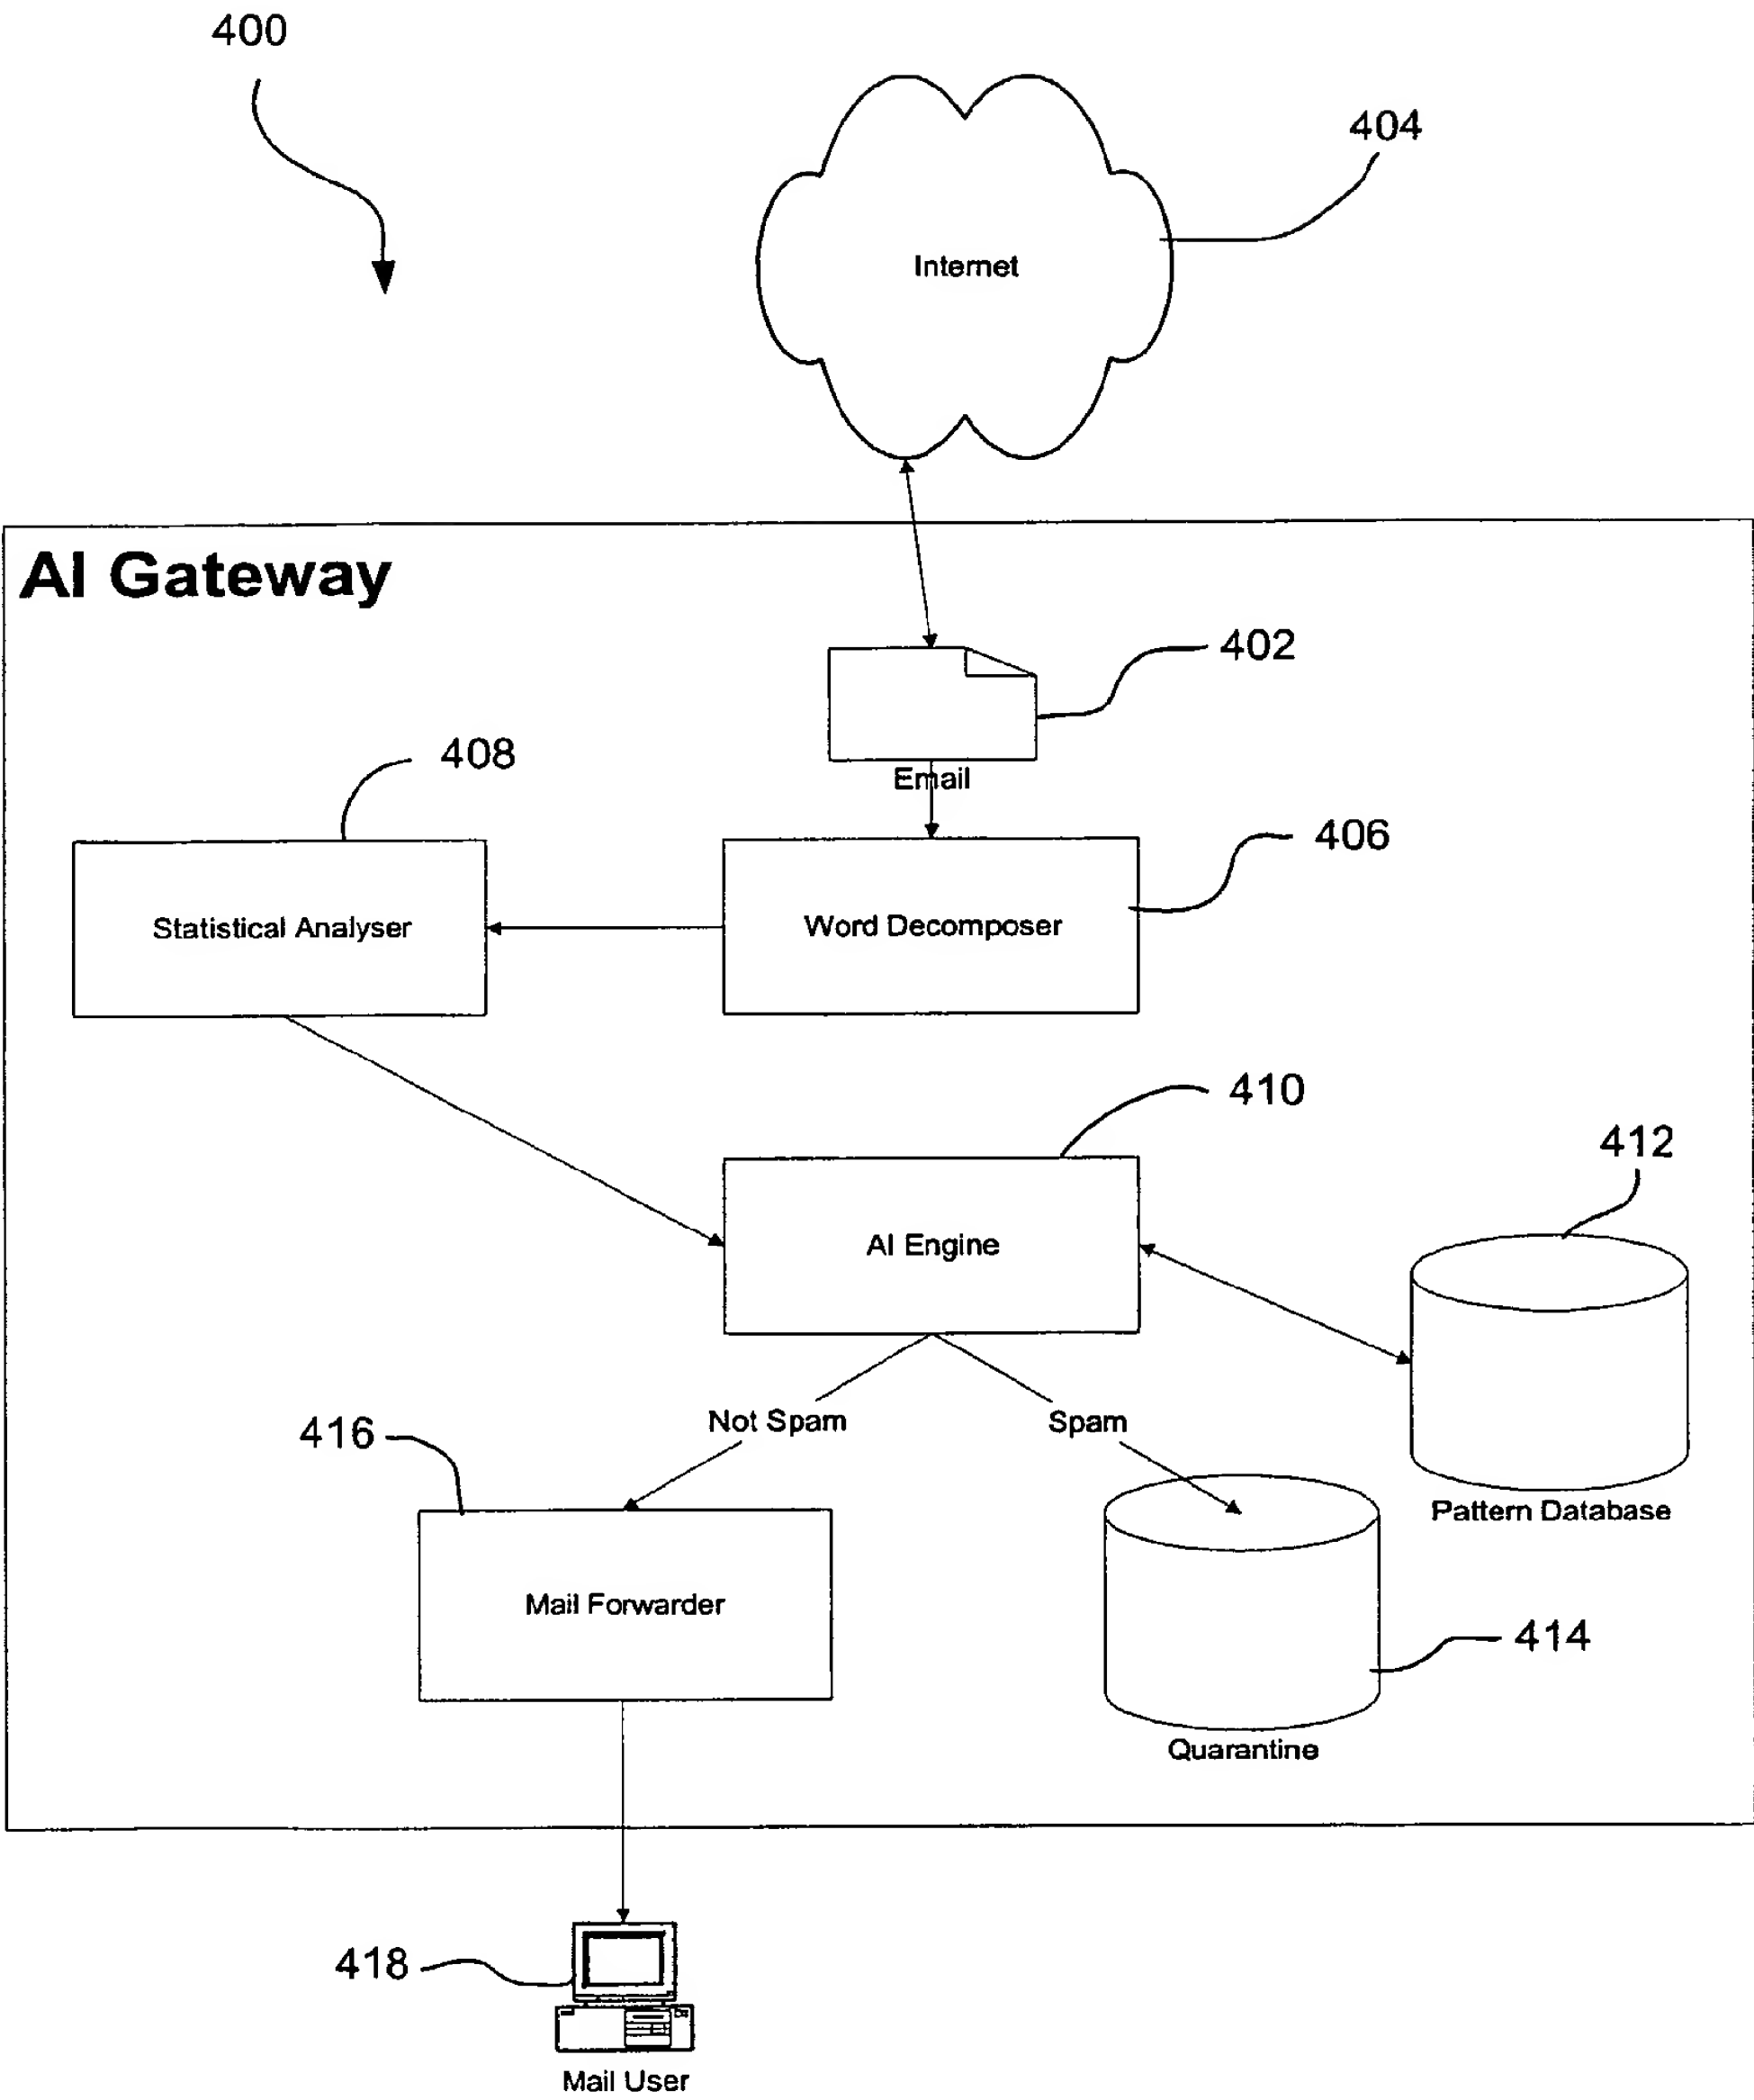
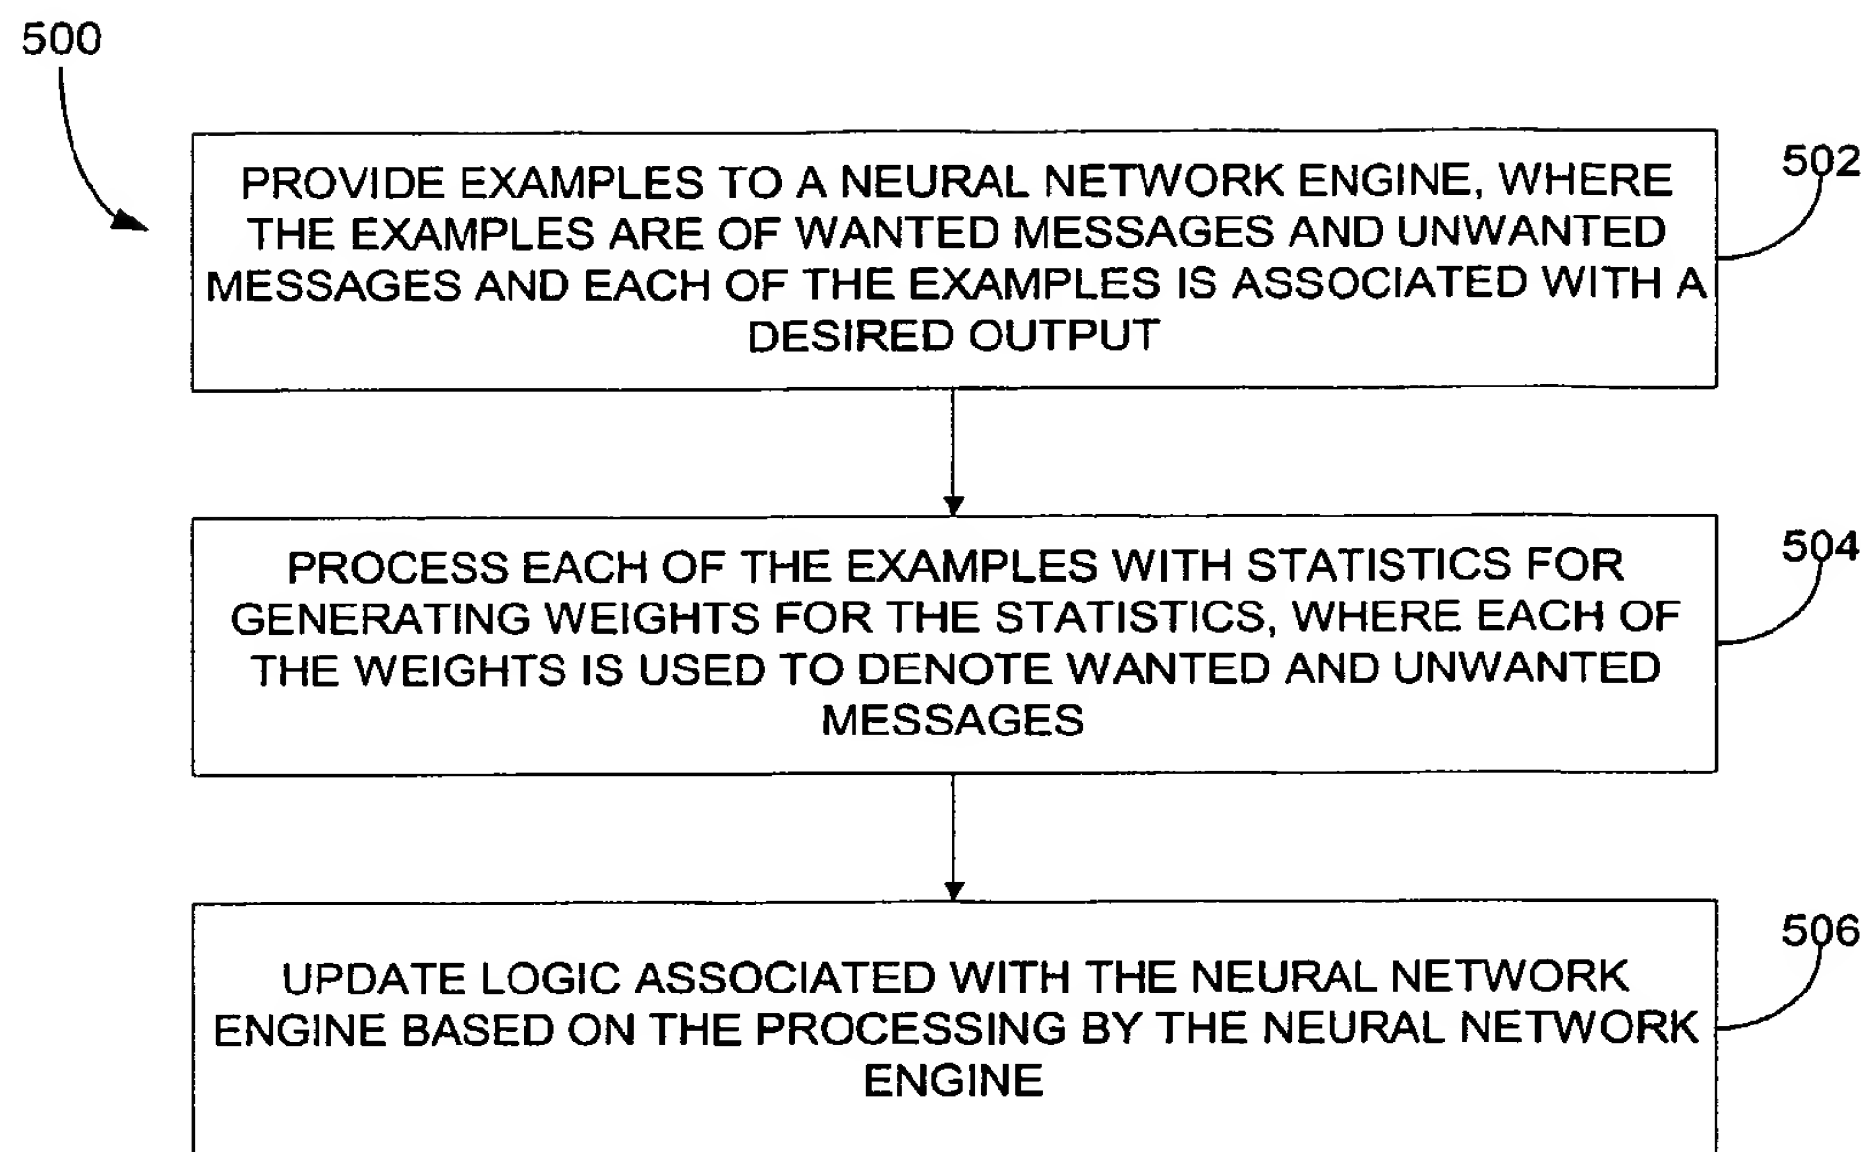


Fig. 4

**Fig. 5**

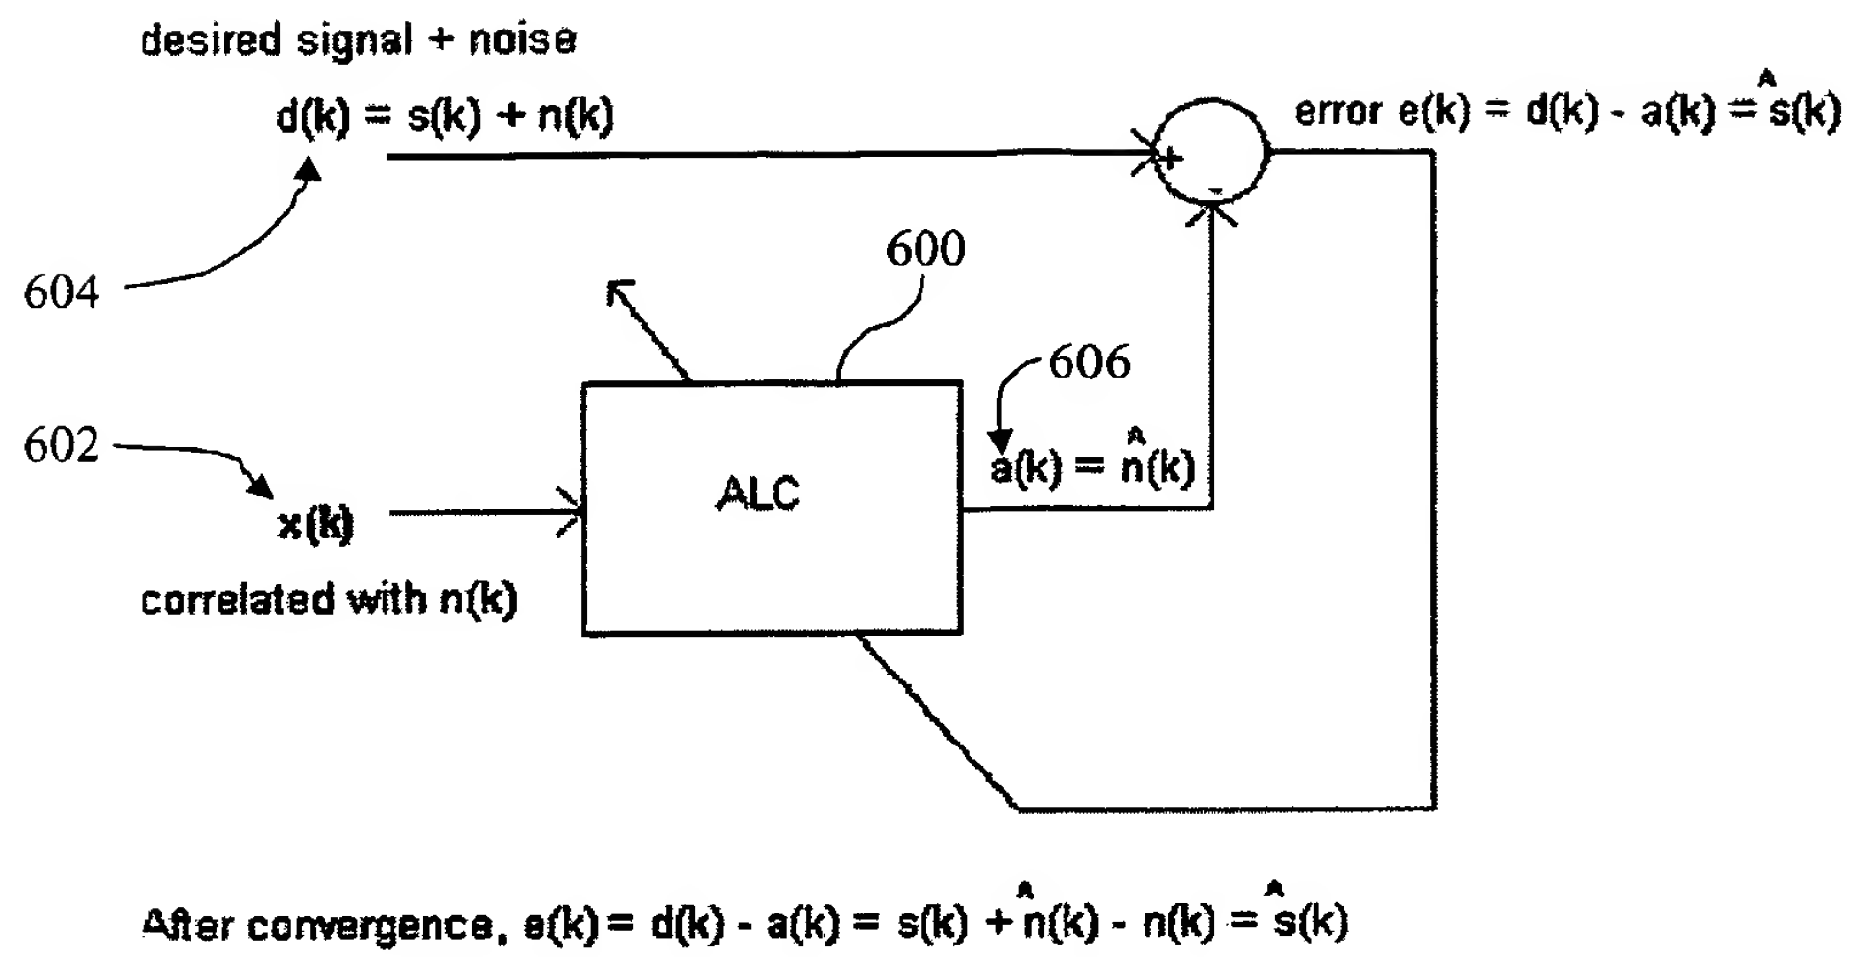


Fig. 6




Fig. 7

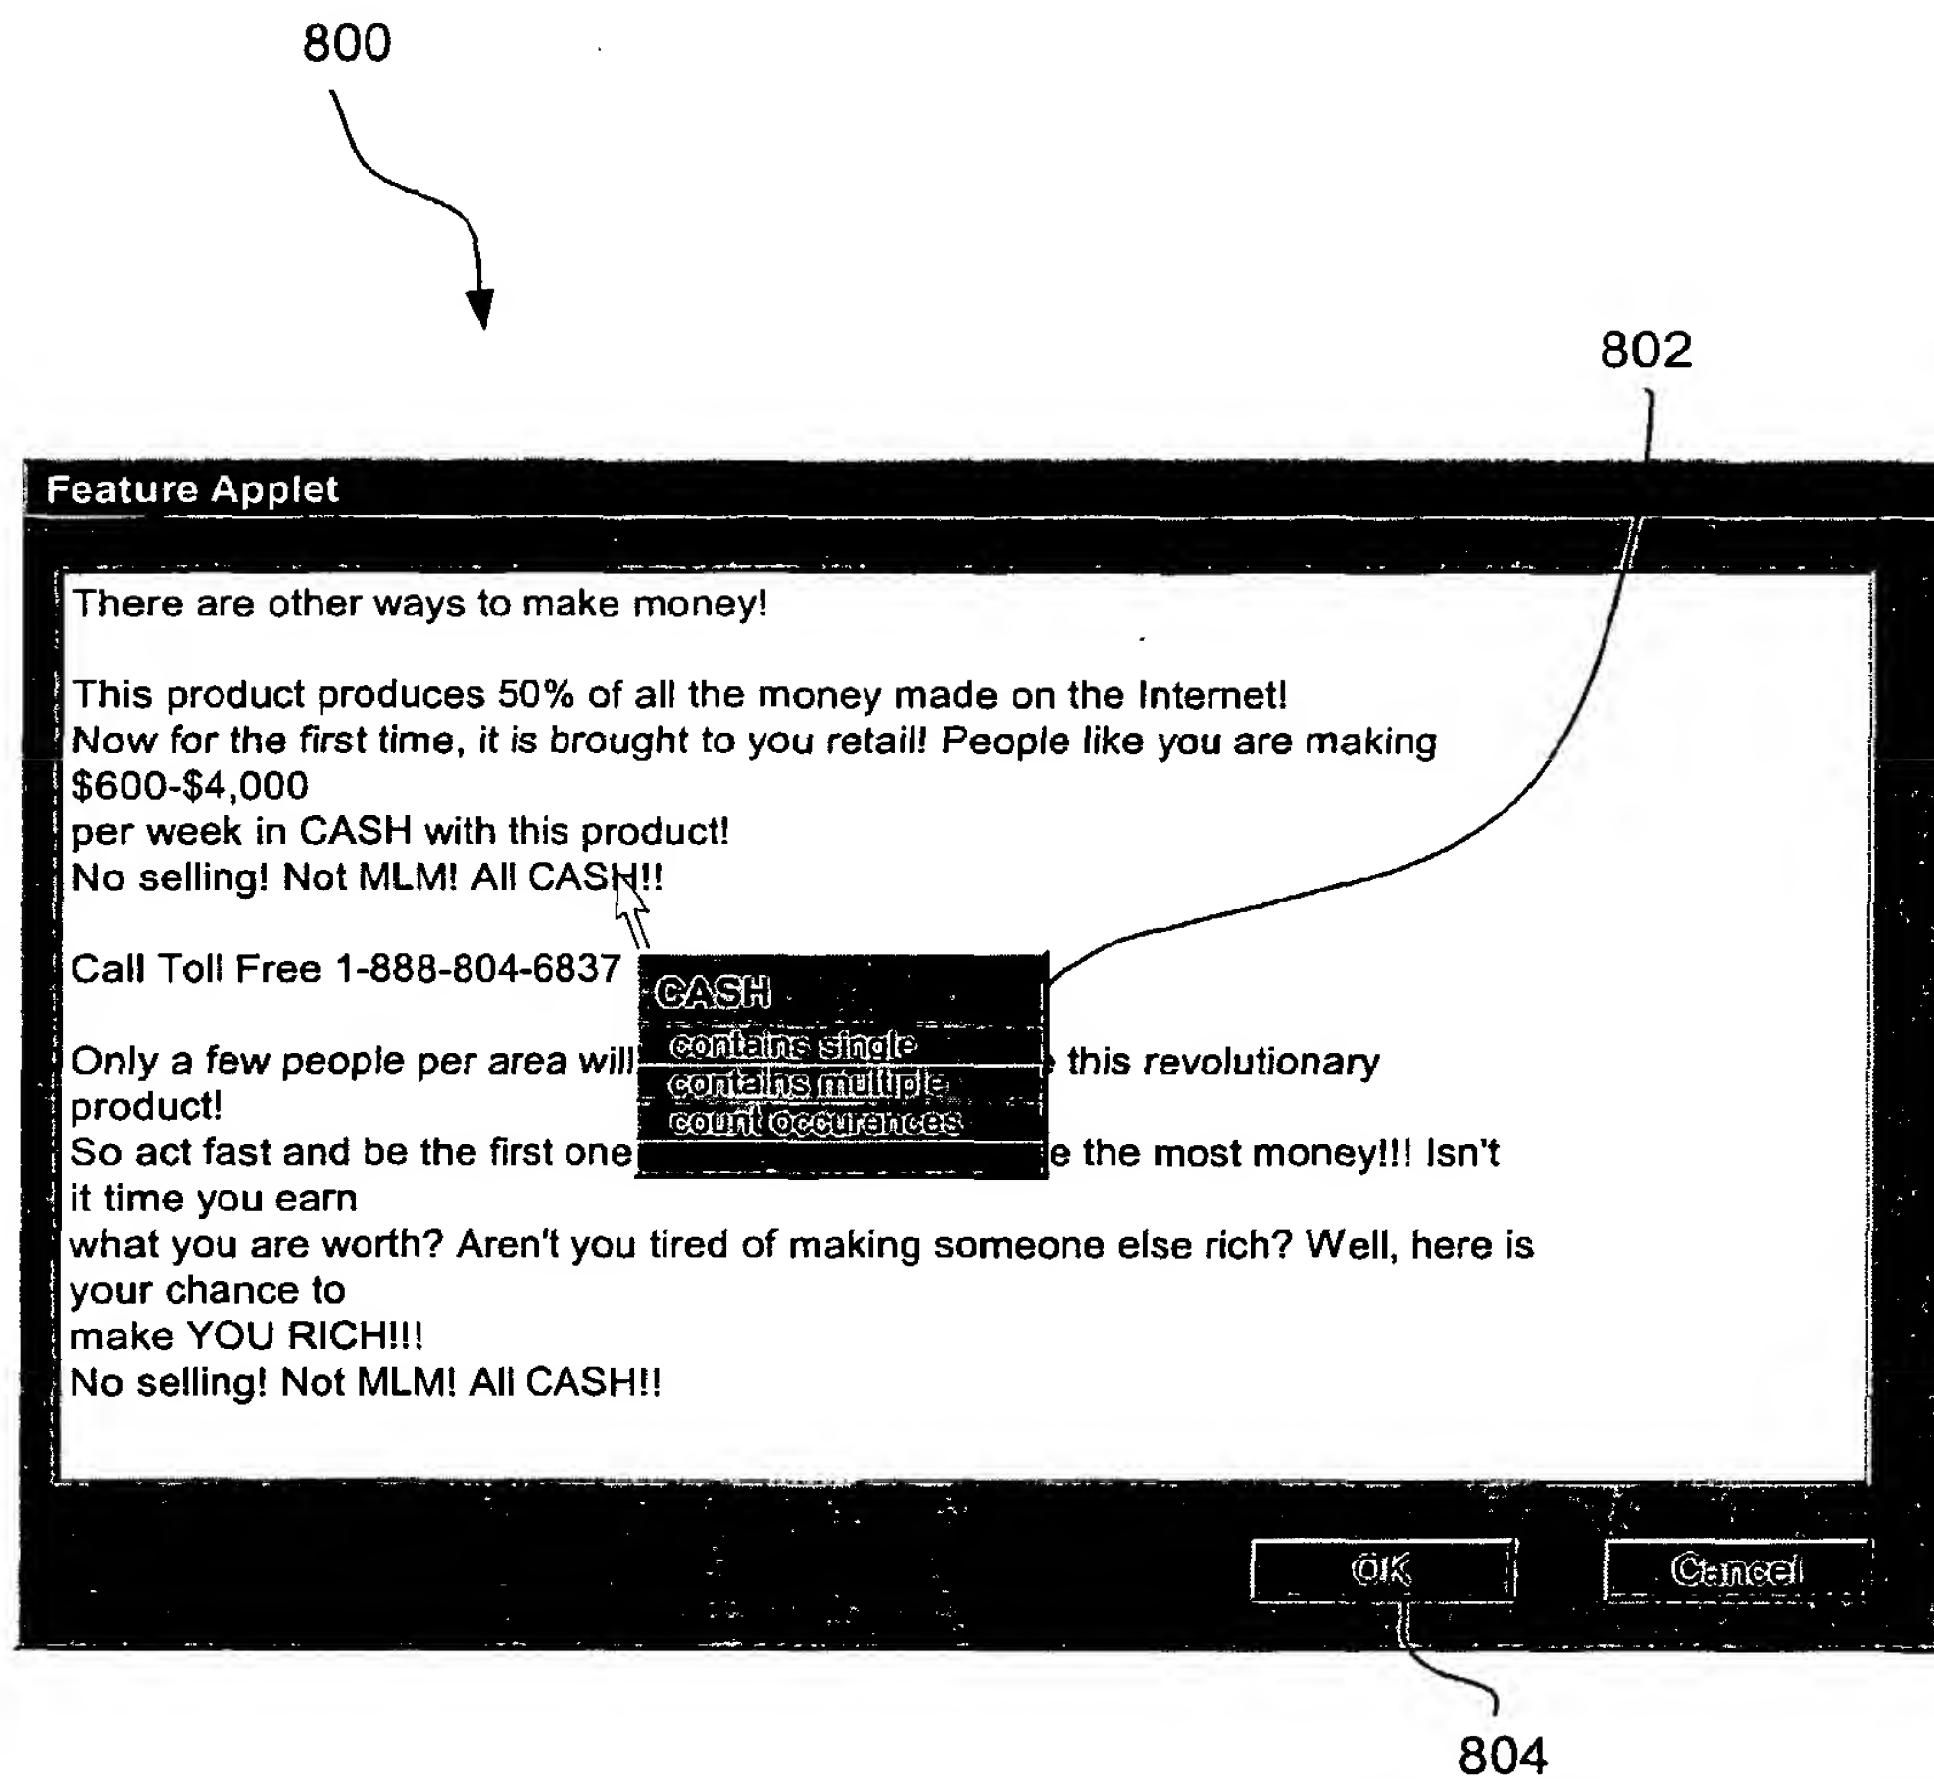


Fig. 8

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US02/23811

A. CLASSIFICATION OF SUBJECT MATTER

IPC(7) : G06F 15/16
US CL : 709/206, 207

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
U.S. : 709/206, 207

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 6,161,130 A (HORVITZ et al) 12 December 2000 (12.12.00), columns 1-5; column 10, lines 65-67; column 13, lines 62-67; column 14, lines 8-11, 19-27, 30-33; column 15, lines 10-53; column 22, lines 25-32	1-17
---		-----
Y		18-21
Y	US 6,330,590 B1 (COTTEN) 11 December 2001 (11.12.01), column 2, lines 20-38; column 3, lines 25-35; column 6, lines 9-13	18-21



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T"

later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X"

document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y"

document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&"

document member of the same patent family

Date of the actual completion of the international search

19 September 2002 (19.09.2002)

Date of mailing of the international search report

29 OCT 2002

Name and mailing address of the ISA/US

Commissioner of Patents and Trademarks
Box PCT
Washington, D.C. 20231

Facsimile No. (703)305-3230

Authorized officer

Ayaz R Sheikh

Telephone No. 703-305-3900

James R. Matthews